arXiv:cs/0101009v1 [cs.PL] 12 Jan 2001

Generation of and Debugging with Logical Pre and Postconditions*

Ángel Herranz Nieva Juan José Moreno Navarro
Universidad Politécnica de Madrid
Departamento LSIIS, Facultad de Informática,
Campus de Montegancedo, Boadilla del Monte,
28660 Madrid, Spain,
mailto:{aherranz,jjmoreno}@fi.upm.es

There are two ways of constructing a reliable system. One is to make it so simple that there are no obvious deficiencies. The other one is to make it so complex that there are no obvious deficiencies.

C.A.R. Hoare

Abstract

This paper shows the debugging facilities provided by the SLAM system. The SLAM system includes i) a specification language that integrates *algebraic specifications* and *model-based specifications* using the object oriented model. Class operations are defined by using rules each of them with logical pre and postconditions but with a functional flavour. ii) A development environment that, among other features, is able to generate readable code in a high level object oriented language. iii) The generated code includes (part of) the pre and postconditions as assertions, that can be automatically checked in the debug mode execution of programs.

We focus on this last aspect. The SLAM language is expressive enough to describe many useful properties and these properties are translated into a Prolog program that is linked (via an adequate interface) with the user program. The debugging execution of the program interacts with the Prolog engine which is responsible for checking properties.

*In M. Ducassé (ed), proceedings of the Fourth International Workshop on Automated Debugging (AADEBUG 2000), August 2000, Munich. COmputer Research Repository (http://www.acm.org/corr/), cs.SE/0101009; whole proceedings: cs.SE/0010035.

1 Introduction

The paper presents the debugging subsystem of the ongoing project SLAM. The SLAM system allows the user to specify a program in a very high level specification language. It is an object oriented formal specification language that integrates *algebraic specifications* and *model-based specifications*. An algebraic specification, like those proposed by the OBJ language family [4, 15, 3] or Larch-LSL [5], is self-contained and independent of the chosen types in the code. Basically a system is described by a decomposition in data types/classes. The behaviour of an operation is described in terms of other operations. On the other hand, specification languages based on abstract models, as Z [19], VDM [9] or Larch interfaces languages [5], have a series of primitive mathematical domains (sets, sequences, tables, etc) to model in an abstract way all data identified during the process of development. Specification of operations uses (explicitly or implicitly) two predicates (pre/postconditions) that describe the relationship between the input and the output by means of a logic formula.

As we will see, SLAM unifies both kind of languages by specifying operations by logical pre and postconditions, but restricting logical formulae to a computable view of quantifiers as traversal operations on data.

The language is the heart of a system that integrates all the stages in the programming process: analysis, design, implementation, documentation and validation.

A SLAM program is not directly stored in a text file. The development environment allows to declare the full specification in a convenient way. Different SLAM components are different views of the system and they can be edited independently. In order to facilitate the understanding of SLAM we will show SLAM elements with a syntax that does not necessarily correspond neither with an internal representation nor the system presentation. They are generated via a friendly interactive interface, so the reader should not pay too much attention to the concrete syntax presented in this paper.

The SLAM system is able to generate executable and readable code in a high level object oriented language. Furthermore, the code contains runtime checkable assertions (i.e. the pre and postconditions of operations) that will help in the declarative debugging process of the program. The user can change the automatically generated code (for instance, to make it more efficient) but the assertions can be reused to ensure that the new code behaves correctly.

The assertions are complex logical formulae, so they are, in principle, difficult to check. The SLAM system generates a Prolog program to check them. The Prolog program is linked with the C++ program to achieve the desired effect.

A similar goal is present in some related works: Anna [11] is an Ada extension that includes facilities for formally specifying the intended behaviour of Ada programs, and the runtime checking of them. The assertions are limited to Ada constructions and they work mainly with algebraic specifications. In [14] a framework for assertion-based debugging is presented: assertions are written using Prolog and a type language so the user cannot specify properties in such an abstract way as in SLAM. Other related works are the *result checking* approach [2] (the user must specify a new algorithm in order to check the

result of the original one), and *relative debugging* [1] (where it is possible to compare the execution of two programs with the same functionality) Disadvantages appear clear, both algorithms/programs might be wrong and some error might be undetectable.

¿From the debugging point of view, the most novel feature of our proposal is that the specification language SLAM is designed as a trade-off between a high expressiveness and the possibility of compilation. It serves both for specifying many useful properties about programs and, at the same time, to generate executable and readable prototypes. It is possible because of the computable behaviour of quantifiers thanks to the object oriented structure of the language.

We present here just a small subset of SLAM features. The discussion of full SLAM is beyond the aims of the paper and can be found in the SLAM webpage [17]. We focus on the most significant SLAM features and how to obtain the Prolog code to check the assertions. The translation from SLAM to Prolog is fully formalized.

The paper organization is the following: we present the language in section 2, compilation schemes into C++ and Prolog in section 3 and the way in which Prolog assertion code is linked with the C++ program are presented in section 4. Section 5 presents conclusions and future work.

2 The Specification Language

SLAM is an object oriented specification language. In this paper we present a functional (i.e. stateless) version of the language. A SLAM program consists in a collection of user defined and predefined classes. A SLAM execution is an expression involving one or more SLAM objects and function (method) calls. The next sections describe the main SLAM components.

2.1 Classes

In SLAM a class is defined by specifying its properties. These properties declare and define all the characteristics of the class. A class is merely declared by giving a class name.

A class can declare *internal attributes*. Attributes specify the internal representation of instances of the class. In SLAM attributes are not defined explicitly as in most of the imperative object oriented languages. A more declarative construction (*algebraic types*) has been directly introduced into SLAM to indicate that a certain syntactical construction belongs to the class.

Every scheme of declaration of attributes defines an internal hidden *attribute constructor*. The attribute constructor and its arguments define an *internal representation* of the object and specification of operations can be written based on that model. For example, if we need a class to model 2D points, we can define the following class:

class Point **is**
  Cartesian (Real,Real) ∈ Point

and now expressions like Cartesian (1,2.5) are instances of the class Point.

There is no restriction on the names of these attribute constructors, it can even be omitted and replaced by a label[1]. For example, next examples show the use of *aggregation* in order to define the classes of segments and stacks:

class Segment **is**
  seg & {source : Point, destination : Point} ∈ Segment

class Stack (Elem) **is**
  contents & {top : **Nat**, elements : **Seq** (Elem)} ∈ Stack

The last example shows how classes can be parameterized. SLAM provides bounded parametric polymorphism in the same style of Pizza [12]. Some predefined SLAM constructions have been used:

- Record type construct with its standard semantics, for instance to define components indicating a source and a destination point.
- Sequences, that are indexed collections of elements. In the Stack class it is used to store the different elements in the stack.

In contrast with other object oriented languages alternative representations are allowed. In SLAM a class can have several alternative sets of attributes, the style is similar to how to define new types in a functional language and quite related to *union types* or *variant records* in some imperative languages. A very similar construct can be found in Pizza [12].

The point example can be extended to represent points in Polar coordinates maintaining the Cartesian representation as another alternative:

class Point **is**
  Cartesian (Real,Real) ∈ Point
  Polar (Real,Real) ∈ Point

Now, the expression Polar (4.5,1.57) also represents an object that belong to the class Point.

The use of alternatives extends to the definition of recursive data structures. Let see a typical example of a recursive data structure in SLAM:

class Tree (Elem) **is**
  Empty ∈ Tree
  Node (Tree (Elem), Elem, Tree (Elem)) ∈ Tree

Up to now, there is no possibility to have references to object, so it is not possible to include side effects. We plan to extend the language in this vein.

2.2 Class relationships

2.2.1 Inheritance

SLAM classes can be defined from the scratch by means of (multiple) *inheritance* from other class. In contrast with some other object oriented languages attribute definitions can be refined by the subclass by using the two usual kinds of inheritance:

[1]This label, expressed as label &, can be omitted when there is just one alternative.

- Extension: New components are added to an attribute constructor.
  class ColouredPoint **is**
    ColouredPoint **isa** Point
    Cartesian (Real,Real, Colour) ∈ ColouredPoint
- Overriding (or refinement): The number of components/arguments in an attribute constructor remains the same but some of the components are replaced by subclasses of the original classes.
  class ColouredSegment **is**
    ColouredSegment **isa** Segment
    seg & {source : ColouredPoint, destination : ColouredPoint} ∈ ColouredSegment

2.2.2 Aggregation

Any class can be specified by declaring a representation based on attributes that are instances of other classes. In the object oriented model the structural relationship between classes is called *aggregation* or *composition*.

2.2.3 Association

Abstract classes understood as *Interfaces* are an important component of SLAM: they can be used to define an abstract behaviour and other classes can indicate that they share this behaviour. The main difference is that this "indication" can be done by two alternative mechanisms. One is *inheritance* and the user need to supply the missing code or to reuse it. The other one is *imitation*. A class The_Class can *imitate* (implement in Java terminology) the behaviour of another abstract class Abstract_Class by matching all its components (attributes if any, functions) by components of The_Class.

Abstract classes and imitation provides the full power of theories (in OBJ terminology) as well as type classes (à la Haskell [8, 10]) playing a more powerful role than C++ templates.

2.2.4 Classes as types

A class defines a datatype. In particular, it is clear that it represents an algebraic datatype when we move labels to constructors and different alternatives are collected in an union type. In the presence of inheritance, type definitions are expanded. Using Haskell [10] notation, the previous classes can be translated to the following types, where constructors names are prefixed by the class name to avoid name clashing:

```
data Point1 = Point1Cartesian (Real, Real)
data Point2 = Point2Cartesian (Real, Real) | Point2Polar (Real, real)
data Segment = SegmentSeg (Point1, Point1)
data Stack a = StackContents (Int, Seq a)  % With an adequate definition of Seq
data Tree a = TreeEmpty | TreeNode (Tree a, a, Tree a)
data ColouredPoint = ColouredPointCartesian (Real, Real, Colour)
data ColouredSegment = ColouredSegmentSeg (ColouredPoint, ColouredPoint)
```

2.3 Predefined classes and relationships

The SLAM system includes a good number of predefined classes as well as some predefined relations between them. Examples of predefined classes are numbers, booleans, characters, dates, angles, strings, records, sequences, lists, dictionaries, etc. An example of the abstract class hierarchy comes from the Collection class. The following figure shows the behavioural organization of these classes, where the layout reflects inheritance:

Collection	
TravCollection	*Traversable collections*
FiniteTravCollection	*Finite traversable collections*
Set	
Multiset	
IndexedCollection	*Indexed collections. They are traversable*
OrderedIdxCollection	*As above, but respecting the order of elements*

2.4 Traversals

Any class that represents a collection of elements could specify a *traversal*. A traversal is either a *simple enumeration* or a sequence of traversals.

- A *simple enumeration* of a class is a triple (first, next, inside), where first is a class function that returns the initial element of the traversal, next is a class function that eliminates the first element and prepares it to give the following element of the traversal, and inside is a class boolean function indicating if the element already belongs to the traversal.

  - Arithmetic ranges [n..m] are objects in SLAM, and can be seen as examples of simple traversals: first returns n, next consist in adding one to the lower limit [n+1..m], and inside is the function that checks if the lower limit is lesser or equal than the upper one $\mathsf{n} \leq \mathsf{m}$.
  - A single element x is also a trivial example of a simple enumeration.
  - A sequence is another example of single enumeration: first returns the first element of the sequence, next eliminates the head of the sequence and inside decides if there are still elements in the sequence.

  As soon as a class has a traversal defined (i.e. is an instance of the TravCollection class) an object of the class can be used in the place of a traversal.

- A general traversal is a non empty sequence of traversal: $[\mathsf{tr}_1, \ldots, \mathsf{tr}_n]$. Traversals in the sequence can be of any class and the order is relevant, because the structure will be traversed using this order. Typically, it contains simple traversals or objects of TravCollection classes. There is just one traversal per class, although several rules to define it can be used depending on the different shapes of the attributes. Traversals are specified in the following way:

  Shape $(\overline{x}) \rightsquigarrow$ Traversal $(\overline{x})$

where Shape $(\overline{x})$ indicates the shape of the object, and Traversal is a simple or general traversal.

An example of a traversal of a binary tree is the element at the root (a simple enumeration), the left subtree (that can be traversed by the same method) and the right subtree.

Empty $\rightsquigarrow$
Node (ls, root, rs) $\rightsquigarrow$ [root, ls, rs]

Notice that changing the order of this definition means a different way to traverse the tree (for instance, first traversing the left subtree, then the root, and finally the right subtree by using the traversal [ls, root, rs]).

A traversal can be interpreted as a way to generate a scheme of code to traverse the data. A simple enumeration can be translated into a single loop or a linear recursion, while a general traversal is moved to a multiple recursive code (even with parallel execution).

2.5 Class Operations

SLAM has a clear functional flavour, so class operations (methods) can be represented by functions. For the shake of simplicity some syntactical distinctions have been introduced to distinguish standard object oriented aspects, although they can be inferred by the system. The classification of operations is the following:

- *Object constructors.* An object constructor is a function designed to create new instances of a class. An object constructor declaration in class Class consists in the name of the constructor[2], and the argument types:
  constructor f : $\overline{T}$ means a function: **function** f : $\overline{T} \rightarrow$ Class

- *Object observers.* Observers allow to access properties of an object without modifying it. An observer of class Class of the form:
  observer f : $(\overline{T}) \rightarrow R$ means a function: **function** f : (Class,$\overline{T}$) $\rightarrow R$

- *Object modifiers.* Modifiers are designed to modify the value of an object. From the functional point of view a modifier of class Class of the form:
  modifier f : $(\overline{T})$ means a function: **function** f : (Class,$\overline{T}$) $\rightarrow$ Class

- *Friend functions.* They are functions that involve two or more objects of the class and they have not special decoration to be declared.

[2]There are no restrictions on constructor names.

Function invocation for constructors, observers and modifiers is done via the usual dot notation, i.e. function f of class Class is called in the object obj by the expression obj.f ($\overline{a}$) or, in case of ambiguity, by obj.Class:f ($\overline{a}$).

In SLAM a function is specified by a set of *rules*. Every rule involves a *precondition* that indicates if the rule can be apply, a *function call scheme*, and a *postcondition* that relates input and output.

The general form of a function specification is the following:

pre :- $P(\overline{x})$
f ($\overline{a}$)
post :- $Q(\overline{x}, \textbf{result})$

where $P(\overline{x})$ is a SLAM formulae involving variables of the arguments ($\overline{a}$) and $Q(\overline{x}, \textbf{result})$ is another SLAM formula. The variable (and reserved word) **result** is always used to represent the computed value of the function. The formal meaning of that specification is the assertion of the following fact:

$$P(\overline{x}) \;\Rightarrow\; Q(\overline{x}, \mathsf{f}(\overline{a}))$$

Although SLAM formulas and expressions will be defined in detail later the reader can accept some examples on class Point:

constructor PointAt : (Real,Real)
pre :- **true**
PointAt (x, y)
post :- **result** = Cartesian (x, y)

observer CoordX : () → Real
pre :- **true**
Cartesian (x,y).CoordX ()
post :- **result** = x

function Distance : (Point,Point) → Real
pre :- **true**
Distance (Cartesian (x1,y1), Cartesian (x1,y2))
post :- **result** = sqrt (sqr (x1 - x2) - sqr (y1-y2))

As in VDM explicit function definitions are allowed. Even more, unconditionally true preconditions can be skipped what provides more concise and readable specifications.

observer CoordY : () → Real
Cartesian (x,y).CoordY () = y

modifier MoveLeft : (Real)
p.MoveLeft (d)
post :- (Distance (**result**, p) = d) **and** (p.CoordY() = **result**.CoordY())

A function specification only indicates the relationship between the result and the arguments, but not necessarily a method to compute the function. Such a method is called a *solution* in SLAM terminology. A solution always gives a *computable expression* to the function (i.e. to the **result** variable in the postcondition). Next section will describe exactly what is a computable SLAM expression.

In those cases where the specification is not a solution, it can be specified in the following way:

function f : $(\overline{T}) \rightarrow$ Class
pre :- $P(\overline{x})$
f $(\overline{a})$
post :- $Q(\overline{x}, \textbf{result})$
$\triangleright$ Solution $(\overline{x})$

where Solution is a computable SLAM expression. For instance, we could have:

modifier MoveLeft : (Real)
p.MoveLeft (d)
post :- (Distance (**result**, p) = d) **and** (p.CoordY() = **result**.CoordY())
$\triangleright$ Cartesian (-p.CoordX (), p.CoordY ())

2.6 SLAM formulae and quantifiers

SLAM formulas are basically logic formulas using the usual logical connectives (**and** - conjunction, **or** - disjunction, **not** - negation, **implies** - implication, and **equiv** - equivalence), predefined and user defined functions and predicates, and *quantifiers*.

SLAM formulas are typed in a similar way than other expressions. In fact, expressions and formulas share the syntax – every formula is an expression of type boolean.

SLAM expressions can combine objects with its own operations. Operations can be combined in any consistent way to produce new expressions. Expressions can use quantifiers of adequate types. Quantifiers are a key feature in SLAM and extend the notion of quantifier in logic. Quantifiers can compute not only the truth of an assertion but also any other value. Moreover quantifiers can be applied to any *collection* object. Examples of collections are sets, sequences, multisets, lists, trees, etc. The basic reading of a quantifier is the traversal of a collection object in order to apply an operation to all the elements in the collection. Therefore, the components of a quantifier are:

- The quantifier itself that generalizes a basic binary operation.
- The description of the collection which is traversed.
- A filter (boolean expression) that indicates which elements in the collection are involved in the quantification. Filters are introduced to avoid the generation of many intermediate objects.
- The quantified expression

A quantifier is written in the following way:

$$\Xi \text{ x } \textbf{in} \text{ D } | \textit{ filter } (\text{x}) \bullet \textit{ expression } (\text{x})$$

where Ξ is the quantifier symbol (that indicates the meaning of the quantifier), D is a collection where the quantifier ranges, x is the quantified variable, *filter* is a boolean expression that restricts the elements of D that are taking into account for the quantification and *expression* is a function to interact with the quantification.

We describe some predefined quantifiers, one for each kind, although SLAM includes all used in logic and many more:

Existential quantification: (**exists** or $\exists$): Decides if there is an element in D where *filter* is true, verifying *expression*. It is a generalization of the logical disjunction.

Universal quantification: (**forall** or $\forall$): Decides if all the elements in C where *filter* is true, verify *expression*. It generalizes the logical conjunction.

Summatory: (**sum** or Σ): Sums all the values *expression (x)* where x ranges in the elements of D where *filter* is true. Obviously, the addition is generalized by this quantifier.

Productory: (**prod** or Π): Multiplies all the values *expression* (x) where x ranges in the elements of D where *filter* is true (by generalizing multiplication).

Counting: (**count** or $\mathcal{N}$): Counts the elements in D where *filter* is true, verifying *expression*.

Selection: (**select** of $\exists^*$): Returns any element in D where *filter* is true, verifying *expression*.

Maximum: (**max** or $\mathcal{M}$): Returns the maximum of all the values *expression* (x) where x ranges in the elements of D where *filter* is true.

Maximizer: (**maxim** or $\mathcal{M}^*$): Returns the value y such that *expression* (y) is the maximum of all the values *expression* (x) where x ranges in the elements of D where *filter* is true.

Minimum: (**min** or μ): Returns the minimum of all the values *expression* (x) where x ranges in the elements of D where *filter* is true.

Minimizer: (**minim** or μ^*): Returns the value y such that *expression* (y) is the minimum of all the values *expression* (x) where x ranges in the elements of D where *filter* is true.

Filter: (**filter** or $\mathcal{F}$): Returns an element of the same class of D but eliminating the elements that verify *expression*.

Map: (**map** or $\mathcal{C}$): Returns an structure with the elements *expression* (x) such that x ranges in the elements of D where *filter* is true.

Sequence constructor: (**seq** or $\mathcal{S}$): Constructs a sequence with indexes in D and elements the values of *expression* (x).

2.6.1 Computable SLAM expression

Any SLAM expression can be computed provided that quantifiers are applied to a FiniteTravCollection class. A FiniteTravCollection is a collection class whose instances can be traversed and they have a finite number of elements.

2.7 Asserting properties

The SLAM compiler is designed to produce executable and readable imperative code. One of the main features of the resulting code is that it contains debugging annotations, similar to the `assert` directive of C. The main idea is to include assertions to check pre and postconditions. However, this is not always possible. SLAM can generate checking code for formulas involving quantifications over objects that can be finitely traversed, although it is not necessary - for postconditions - that they assign a value to **result**.

When a formula cannot be checked completely, the user (maybe later with the help of the SLAM compiler) can indicate three possible situations. For simplicity we apply them to the postcondicion but the same annotations can be done for the precondition.

- The full formula can be checked. This is done by adding the keyword **check** to the formula. This is the case when the specification is yet a solution and in this case no annotation is needed.
  function f : $(\overline{T}) \rightarrow$ Class
  pre :- $P(\overline{x})$
  f $(\overline{a})$
  post :- **check** $Q(\overline{x}, \textbf{result})$

- Just a part of the formula can be checked. This means that it is a part of a conjunction, so we assume that first part of the formula cannot be checked, and the second one is marked by the keyword **and_check**.
  function f : $(\overline{T}) \rightarrow$ Class
  pre :- $P(\overline{x})$
  f $(\overline{a})$
  post :- non-checkable-property($\overline{x}$, **result**) **and_check** checkable-property($\overline{x}$, **result**)

- The formula cannot be checked at all or the part that can be checked is not significant. In this case, the user can provide an approximation of the property that is checkable by placing it after the keyword **either_check**.
  function f : $(\overline{T}) \rightarrow$ Class
  pre :- $P(\overline{x})$
  f $(\overline{a})$
  post :- non-checkable-postcondition($\overline{x}$, **result**) **either_check** checkable-approximation($\overline{x}$, **result**)

2.8 An additional example

Let us show a more complete but simple example of a SLAM specification. The problem is to handle transactions between saving banks. The input is a collection of transactions (with the source bank, the destination bank and the amount) and the collection of bank names. The function to specify must compute the final profit/debit for each bank.

```
class Transaction is
  public tran & {source : String, destination : String, amount : Real} ∈ Transaction
class Bank is
  public bank & {name : String, amount : Real} ∈ Bank
constructor MakeBank : (String,Real)
MakeBank (n, a) = {n, a}
class CTransaction is
  ctran & Seq (Transaction) ∈ CTransaction
observer FinalAmount : Seq (String) → Seq (Bank)
pre :- banks ≠ []
FinalAmount (ctrans, banks)
post :- check  length (result) = length (banks) and
               forall i in [1..length(result)] •
                 result(i).name = banks(i) and
                 result(i).amount =  (sum t in ctrans | t.source = banks(i) • t.amount) -
                                     (sum t in ctrans | t.destination = banks(i) • t.amount)

▷  map n in banks • MakeBank (n,  (sum t in ctrans | t.source = n • t.amount) -
                                  (sum t in ctrans | t.destination = n • t.amount))
```

2.9 Additional features

SLAM includes some other features that are quite convenient for program specification. However, due to the lack of space we have omitted those characteristics that are less interesting for the purpose of this paper. Among them we can mention: *Selectors*, that extend patterns for hidden constructors of classes, *Internal laws*, to simplify the data stored in a class after any operation, a broad notion of visibility in case of aggregation, etc.

3 Compiling SLAM programs

3.1 Compiling SLAM to an imperative object oriented language

One of the main features of the SLAM is that it is able to generate code from a specification. We will use C++ ([18]) as the target language. However the ideas are applicable to any language with similar characteristics (Java, Smalltalk, etc.). In this section we only sketch the main ideas for the compilation because this is not the main goal of the paper. The basis for the compilation are:

- SLAM classes are translated to C++ classes.

  An SLAM class defines a declarative type via its attributes. This type is translated into an imperative type. The C++ class contains a single attribute of this type. The construction of an imperative type from a declarative one is a key feature for the compilation. Constructors are translated to records, alternative definitions are moved to union types while recursive definitions need the inclusion of a pointer.

- SLAM class relationships are almost mimic by C++ class relationships.

  - Aggregation is already included in object oriented languages.
  - Most of the power of inheritance on SLAM is included in C++, except the inheritance modification of the state. Attribute constructor extension (or record extension) can be modelled by C++ attribute extension (i.e. additional attributes can be supplied in the subclass), while state overriding can be implemented by attribute redefinition[3].
  - Association can be simulated by inheritance and/or the use of templates.

- SLAM predefined classes are implemented by specifically designed C++ classes, although some of them need a special treatment by the compiler.

- SLAM functions are implemented by C++ methods in traversable classes. The SLAM classification of operations are easy to map into C++ classification.

- Traversals are implemented by including the operations `first`, `next`, and `inside` as class methods. Quantifiers are implemented by adequate loops.

- An additional operation `serialize` is included in all the classes to store an element into a file using the XML format. This operation is also automatically generated.

In our previous example, we can obtain automatically a C++ code similar to the following one[4]:

[3] In C++, when a subclass declares an attribute with the same name of an attribute of the superclass, this last one is hidden by the newest.

[4] Some parts have been skipped, and some simplifications like using arrays as the **Seq** implementation have been included to improve readability.

```
// in file transaction.h
class Transaction {
 public:
   String source;
   String destination;
   float amount;

   void serialize ();
};
```

```
// in file bank.h
class Bank {
  public:
    String name;
    float  amount;

    void Bank (String n, float a) {
      name = n;
      amount = a;
    }
    void serialize ();
};
```

```
// in file ctransaction.h
class CTransaction {
  public:
    void FinalAmount (
      String banks [MaxBanks],
      Bank result [MaxBanks]);
    void serialize ();

    Transaction trans [MaxTrans];
};
```

```
// in file ctransaction.cpp
void CTransaction::serialize () {
...
}

void CTransaction::FinalAmount (
  String banks [MaxBanks],
  Bank result [MaxBanks])
{
  pre-check (FinalAmount (banks, MaxBanks));
  for (int i = 0; i <= MaxBanks; i++){
    float amount = 0;
    for (int j = 0; j <= MaxTrans; j++){
      if (trans[j].source == banks [i])
        amount =+ trans[i].amount;
      else if (trans[j].destination == banks[i])
        amount =- trans[i].amount;
    }
    result[i] = new Bank (banks [i], amount);
  }
  post-check (FinalAmount
                (banks, MaxBanks, result))
}
```

This automatically generated code is not necessarily intended to be an efficient code. In fact, it could be considered a prototype code, although in many cases the code is good enough. In those situations, or even in situations in which SLAM cannot generate code, the user can modify it and write an optimizing code.

In our previous example, the programmer can decide that it is better to traverse first the transaction collection and then search the name of the banks to update the result. The new code is:

```
void CTransaction::FinalAmount (String banks [MaxBanks],
                                Bank result [MaxBanks]) {
  pre-check (FinalAmount (banks, MaxBanks));
  for (int i = 0; i <= MaxBanks; i++){
    result[i].name = banks [i];
    result[i].amount = 0;
  }
```

```
  for (int j = 0; j <= MaxTrans; j++){
    for (i = 0; i < MaxBanks; i++) {
      if (trans[j].source == result [i].name)
        result[i].amount =+ trans[i].amount;
    }
    for (i = 0; i < MaxBanks; i++) {
      if (trans[j].destination == result [i].name)
        result[i].amount =+ trans[i].amount;  // There is a bug here
    }
  }
  post-check (FinalAmount (banks, MaxBanks, result))
}
```

Notice that the code contains a bug when the destination bank is found. Instead of decreasing the amount it is increased. Program assertions remain valid, so the bug will be detected when the program is executed in the debug mode.

3.2 Compiling SLAM to Prolog

More relevant for the purpose of the paper is the compilation of SLAM programs and assertions into a Prolog piece of code.

We use the Ciao/Prolog system [6] as the target language. The main reasons are our knowledge of the system and our easy access to the Ciao/Prolog system developers, the availability of some higher order features that simplify the implementation of quantifiers, and the existence of an interface to connect C/C++ and Prolog programs. Anyway, the most important ideas presented here are independent of the Ciao/Prolog system and are applicable to any other Prolog system.

We have already explained that the semantics of classes are algebraic types. Instances of classes will be translated using Prolog functors as data constructors. This is the usual way in which other declarative languages translates algebraic types into Prolog. The name of the class is added to the name of the constructor in order to avoid clashing in the presence of inheritance.

A formal specification of the translation of function definitions is given in figure 1. Every function is compiled into three predicates that implements different functionalities. For every function f with arity n we will have the predicate `'sol-`f`'` that implements the function f. The arity of this predicate is $n + 1$ where n is the arity of the function (remember that member functions in SLAM are functions with an element of the class as the first argument). Predicates `'pre-`f`'` and `'post-`f`'` will implement the *checking* predicates of the pre and postconditions of the function f.

The complete translation function *translate* (see figure 1) accepts a SLAM program as argument and returns a Prolog program. The definition uses several auxiliary functions.

- *trans_rule* accepts a rule and the class where the rule appears and returns three Prolog clauses: for `'sol-`f`'`, `'pre-`f`'` and `'post-`f`'`. In the definition a rule is represented by a record with six fields: *class* (the class where it is defined), *fname* (function

$$
\begin{aligned}
\mathit{translate}(SP) = {} & \textstyle\bigcup_{\mathit{rule}\,\in\,SP} \mathit{trans_rule}(\mathit{rule}, \mathit{rule.class}) \cup \\
& \textstyle\bigcup_{\mathit{miss}(f,C)} \mathit{trans_inh}(f, C, \mathit{defines}(f, C)) \cup \\
& \textstyle\bigcup_{C\,\in\,SP} \mathit{trans_read}(C) \cup \\
& \textstyle\bigcup_{C\,\in\,SP} \mathit{trans_trav}(C) \cup \\
& \mathit{trans_predefined}
\end{aligned}
$$

$$\mathit{trans_rule}(r, C) = \mathit{trans_pre}(r, C) \ \cup\ \mathit{trans_post}(r, C) \ \cup\ \mathit{trans_sol}(r, C)$$

$$\mathit{trans_pre}(r, C) = \texttt{'pre-}f\texttt{'}(\mathit{trans_term}(C, r.\mathit{args}))\ \texttt{:-}\ \mathit{trans_expr}(r.\mathit{pre}, \texttt{Pre}, C)\texttt{, Pre == true.}$$

$$\mathit{trans_post}(r, C) = \texttt{'post-}f\texttt{'}(\mathit{trans_term}(C, r.\mathit{args}), \texttt{Result})\ \texttt{:-}\ \mathit{trans_expr}(r.\mathit{post}, \texttt{Post}, C)\texttt{,}\ \texttt{Post == true.}$$

$$\mathit{trans_sol}(r, C) = \texttt{'sol-}f\texttt{'}(\mathit{trans_term}(C, r.\mathit{args}), \texttt{Result})\ \texttt{:-}\ \mathit{trans_expr}(r.\mathit{sol}, \texttt{Result}, C)\texttt{.}$$

$$\mathit{trans_inh}(f, C, C') = \textstyle\bigcup_{a\in\mathrm{C}}\ \texttt{'sol-f'(}\mathit{trans_term}(C', a)\texttt{, R) :-'sol-}f\texttt{'(}\mathit{trans_term}(C', a)\texttt{, R).}$$

$$\mathit{trans_expr}(\texttt{c (}e_1\texttt{, ..., }e_n\texttt{)}, \texttt{V}, C) = \mathit{trans_expr}(e_1, \texttt{X}_1, C)\texttt{, ...,}\mathit{trans_expr}\ (e_n, \texttt{X}_n, C)\texttt{,}\ \texttt{V is}\ \mathit{trans_term}(C, c)\texttt{(X}_1\texttt{, ..., X}_n\texttt{).}$$

$$\mathit{trans_expr}(\texttt{f (}e_1, \ldots, e_n\texttt{)}, \texttt{V}, C) = \mathit{trans_expr}(e_1, \texttt{X}_1, C)\texttt{, ...,}\mathit{trans_expr}\ (e_n, \texttt{X}_n, C)\texttt{,}\ \texttt{'sol-}f\texttt{'(X}_1\texttt{, ..., X}_n\texttt{, V).}$$

$$\mathit{trans_expr}(\texttt{obj.f (}e_1, \ldots, e_n\texttt{)}, \texttt{V}, C) = \mathit{trans_expr}(\texttt{f (obj, }e_1, \ldots, e_n\texttt{)}, \texttt{V}, C)$$

$$\mathit{trans_expr}(\texttt{C':f (}e_1, \ldots, e_n\texttt{)}, \texttt{V}, C) = \mathit{trans_expr}(\texttt{f (}e_1, \ldots, e_n\texttt{)}, \texttt{V}, C')$$

$$\mathit{trans_expr}(\textbf{result}, \texttt{V}, C) = \texttt{V is Result.}$$

$$\mathit{trans_expr}(e_1\ \texttt{log_op}\ e_2, \texttt{V}, C) = \mathit{trans_expr}(e_1, \texttt{X}_1, C)\texttt{,}\mathit{trans_expr}(e_2, \texttt{X}_2, C)\texttt{,}\ \texttt{log_op (X}_1\texttt{, X}_2\texttt{, V).}$$

$$\mathit{trans_expr}(\Xi\ x\ \textbf{in}\ \mathit{coll} \mid \mathit{filt} \bullet e, \texttt{V}, C) = \texttt{'quan-}\Xi\texttt{'(}\mathit{coll}\texttt{, }\mathit{filt}\texttt{, }e\texttt{, V)}$$

Figure 1: Formalization of the translation.

name), *args* (arguments), *sol* (the solution proposed by the rule), *pre* (checkable part of the precondition), and *post* (checkable part of the postcondition).

- *trans_sol*, *trans_pre*, and *trans_post* are responsible for producing clauses for the function implementation and for checking the pre and postconditions.
- *trans_expr* produces Prolog atoms in order to translate a SLAM expression. *trans_expr*(e, X, C) ensures that the Prolog variable X stores the computation of the expression e in the class C.
- *trans_inh* is designed to produce the clause associated to an inherited function f definition in class C checked by *miss*. *defines* returns the ancestor class where a function is declared (see below for further details).
- *trans_term* translates a term in a class into the corresponding Prolog representation (following the name policy described above).

In the compilation scheme presented in figure 1 c represents an attribute constructor while f represents a function. $a \in C$ *produces* all possible skeletons of terms in class C.

log_op is any SLAM logical operator computed in Prolog by the corresponding `log_op` predicate[5]. *trans_predefined* contains a Prolog implementation of SLAM predefined operations. Anyway the scheme deserves a deep discussion on some aspects:

It is important to note that the same predicate name is used to compute the same function in all the classes coming from the ancestor class. Some few words are needed to justify this decision and to show that it really handles the overloading produced by the use of inheritance. As Prolog has no types, a predicate can be applied to any number of functors on the top of arguments. Remember that attribute constructors in different classes have different names. If f is a member function of C and the functions is applied to an expression e the generated Prolog atoms are:

$$\textit{trans_expr}(x, \texttt{Y}, C)\texttt{, 'sol-}f\texttt{'(Y,R)}$$

Notice that the clauses of `'sol-`f`'` includes all the definitions of f in the complete class hierarchy, so it does not matter the exact subclass A of C which x belongs to. As `Y` will not contain free variables and the attribute constructors in all the classes have different names, only the adequate original rule is applicable. A similar approach is used for `'post-`f`'` and `'pre-`f`'`.

Patterns are translated into Prolog patterns, so parameter passing is handle by unification. This use of unification is quite efficient in Prolog because our programs only use the *read mode* of the underlying WAM machine. Function composition is handled by flattening the functional structure and using auxiliary variables to connect the predicates.

There are additional technical details that must be solved. One appears when a function in class A is fully inherited from a class C, i.e. there are no rules for f in A. As the attribute constructors in both classes may be different, we need to add a clause to transforms the arguments to have the constructors in C and then call recursively the same predicate. Suppose that the attribute constructor in C is `c` and the attribute constructor in A is `c'` with possibly more attributes. The generated clause is:

$$\texttt{'sol-}f\texttt{'(c'(}\overline{x}\texttt{,}\overline{y}\texttt{), R) :- 'sol-}f\texttt{'(c(}\overline{x}\texttt{), R).}$$

The second one is how to compile a function call with explicit indication of the class. This expression could involve a type casting from an argument to an ancestor class. For each class C we could generate a predicate `'to-`C`'` that transforms all the elements of any successor class A into C. A call of the form $\mathsf{obj}.C{:}\mathsf{f}\ (\overline{x})$ is translated into the atoms

$$\texttt{'to-}C\texttt{'(obj,Obj), 'sol-}f\texttt{'(Obj,}\overline{x}\texttt{,R)}$$

Notice that the dispatch table used for the compilation of inheritance is simulated here by the use of different functors names (keys into the table) as well as the unique predicate with all the possible implementations of the function (the table itself). One can argue that in case of an intensive use of inheritance this huge predicate can be a source of inefficiency. Remember that the Prolog program is just used for debugging, so efficiency is not a crucial point. On the other hand, this possible source of inefficiency is partially solved by the indexing included in WAM-based Prolog implementations.

[5]Notice that they cannot be implemented by the similar Prolog operations (conjunction $\leftrightarrow$ `,`, etc.) because SLAM boolean functions compute `false` as a valid result. We leave their definitions as a trivial exercise.

Translation of traversals and quantified expressions follow an interesting approach. The *trans_trav* translation scheme is applied for any class C implementing the TravCollection and consists in the generation of a predicate `in`. The first argument of the predicate is the traversable object, the second argument *produces* by backtracking all the values in the collection. A simple traversal is translated in this way:

```
in(O,X) :- first(O,X).
in(O,X) :- next(O,O2), inside(O2), in(O2,X).
```

while a traversal definition using a sequence of traversals is translated into several clauses, one for each traversal in the sequence:

```
in(<<shape>>,X) :- in(<<tr_1>>,X).
...
in(<<shape>>,X) :- in(<<tr_n>>,X).
```

A trivial translation of a quantifier definition is got by using the following predicate quantifier. Ciao/Prolog allows a special syntax `f (`$\overline{\mathsf{A}}$`) := `*Expr* `:-` *condition* to compute the function `f` as the expression *Expr* provided that *condition* holds. The Ciao/Prolog higher order contains the usual higher order functions, like `~foldr(`$f, s, [x_1, \ldots, x_n]$`)` that computes the value of $f(x_1, f(x_2, \cdots f(x_n, s) \cdots)$. There is also a special syntax for lambda expressions (called predicate abstractions) `\(`$\overline{\mathsf{X}}$`, R) :- Expr`, where `R` stores the computed expression `Expr` applied to the arguments $\overline{\mathsf{X}}$. We show the implementation of some quantifiers as an example.

```
quantifier(Binop,Seed,Collection,Filter,Expression) := ~foldr(Binop,Seed,L) :-
     findall(Y,(in(Collection,X),Filter(X),Expression(X,Y)),L).
'quan-all'(Collection,Filter,Expression, R) :-
     R = quantifier(log_and,true,Collection,Filter,Expression).
'quan-sum'(Collection,Filter,Expression, R) :-
     R = quantifier(+,0,Collection,Filter,Expression).
'quan-map'(Collection,Filter,Expression, R) :-
     R = quantifier(\(X, L, R) :- R = [X|L],[],Collection,Filter,Expression).
```

In some cases the filter predicate can generate values in a very efficient way and a more efficient translation could be `(Filter(X), in(Collection,X))` in the condition of the "findall´´ construct. In principle, we have not information for choosing between both alternatives, unless we use some techniques borrowed from program analysis.

It is worth mentioning that, by this translation, it will be clear that the SLAM selection quantifier provides the full power of Prolog search, what enhances the expressive power of SLAM.

Notice that the checkable part of function pre and postconditions are translated also into special predicates. The only tricky question is the use of the same argument patterns of the function instead of the variables to ensure the adequate parameter passing even for the conditions.

The last part of the translation that must be explained is the goal of the *trans_read* (C) clauses. For each class C a predicate with name `'read-`C`'` is defined in order to read an element of class C from a class. The format of the data in the file is exactly the same that

is stored by the counterpart operation (`serialize`) of the C++ class for C. Due to the lack of space, we omit the definition although it is an easy exercise.

Let us show the Prolog code for our example. Again we have omitted some parts, and included some simplifications to make it more readable, like the implementation of sequences as lists or the inclusion of intermediate predicates to compute the quantified expressions.

```
amount(tran(S,D,A),A).
'sol-MakeBank'(N, S, bank(N,S)).
sumSource(Ctr,N,V) :- 'quan-sum'(CTr,(\:(tran(S,D),R) :- S == N -> R = true; R = false), amount
sumDest(Ctr,N,V) :- 'quan-sum'(CTr,(\:(tran(S,D),R) :- D == N -> R = true; R = false), amount,
'sol-FinalAmount'(CTr, Banks, Result) :-
      'quan-map'(Banks, true, \(N, R) :-
                                sumSource(CTr, N, A1), sumDest(CTr, N, A2), A is A1 - A2,
                                R = 'sol-MakeBank'(N,A)).
'pre-FinalAmount'(CTr, Banks) :- Length(Banks, L), L > 0.
'post-FinalAmount'(CTr, Banks, Result, Post) :-
      Length(Banks, L1), Length(Result, L2), L1 == L2,
      all(1..L1, true, (\(I, R) :- nth (Result, I, bank (N, A)), nth (Banks, I, N),
                                   sumSource (CTr, N, A1), sumDest (CTr, N, A2),
                                   A == A1 - A2), true).
in([X|_], X).
in([_|L], X) :- in(L, X).
in(N ..  M,N).
in(N..M,R):- N <= M, NN = N+1, in(NN..M, R).
```

It is obvious that the transformation can be improved avoiding several sources of inefficiency: generation of useless intermediate variables, predicate calls that can be folded, construction of the result in the head of the clause, etc. However, it is worth mentioning that the Ciao/Prolog system includes several optimizations based on partial evaluation and global analyses in the compilation process that can solve part of the inefficiencies. The system can eliminate part of the conditions to be checked ([13, 7]).

4 Execution of assertions

We have developed a C++ library called `"check.h"` including the macros needed to execute the assertions. Basically, the library contains three operations: `pre-check` (*function name, arguments*), to check preconditions, `post-check`(*function name, arguments, result*), to check postconditions, and `post-return-check`(*function name, arguments, result*), to check a postcondition and return a value.

The interface between C/C++ and Prolog is very primitive. A string containing the goal is used to call the Prolog program. The arguments (and the result, in the case of the postconditions) are stored in a file ***file*** via the `serialize` operation. A process is spawned to run the Prolog program, which reads from the file the arguments/result and calls the corresponding predicate. The pseudo-code for `post-return-check` (*fname,* $\overline{args}$*, result*) is the following:

```
r = result;
fd = fopen (file);
arg1.serialize (fd);
⋮,
argn.serialize (fd);
r.serialize (fd);
fclose (fd);
fork (prolog ('post-fname');
assert (return code of process);
return r;
```

The other cases are analogous. We only need to add the Prolog program two clauses for each function f of the form:

```
'pre-f' :- 'read-C1'(A1), ..., 'read-Cn'(An),'pre-f'(A1, ..., An).
'post-f' :- 'read-C1'(A1), ..., 'read-Cn'(An),'read-Cr'(Result),
            'post-f'(A1, ..., An, Result).
```

The use of the file for exchange information between the C++ program and the Prolog program can be optimized using a pipe to link the processes. As a pipe is handled as a file, there is no significant difference w.r.t the previous scheme.

There is an alternative possibility. First of all, we modify `serialize` to construct a Prolog representation of data as a string. Then a Prolog goal of the form:

```
assert (arguments (arg1.serialize (), ...argn.serialize (), result.serialize ()),
arguments (A1, ..., An, Result), 'post-f' (A1, ..., An, Result).
```

is used to call the Prolog engine.

The Prolog *assert* is needed to ensure that the arguments are compiled. The main problem with this approach is that when the data are large, the string representation of the goal is longer than the accepted size by the Prolog compiler.

Additionally, the debugging execution produces a report on the pre/postconditions checked, indicating those conditions that where marked as a subformula (**and_check**) or alternative condition (**either_check**) of the full pre/postcondition.

5 Future work and Conclusion

We have presented the validation facilities of the SLAM system. From a high level specification the system generates code to execute the system including assertions to check pre and postconditions of the program operations. These high level logical formulae are compiled into Prolog and the resulting program is responsible for checking them. The compiled SLAM program is linked with the assertion Prolog program to achieve the debugging facilities that are specially useful when the programmer decides to modify the automatically generated code.

The overall efficiency is acceptable. The execution of the program in the debug mode needs between 3-4 times the execution of the program itself.

The SLAM project is at the early stages of its development. The specification language is not fully defined yet and the system (environment, compilers, libraries, etc.) is under development. We plan to have a prototype with the debugging facilities by the date of the workshop.

The SLAM project and its debugging facilities seem to be a very useful tool to develop high quality programs, i.e. error free with respect to the specification, clean, easy to read and manipulate to achieve modifications either in the specification or in the generated code, fully documented and including high level declarative debugging facilities to allow optimizations in a reliable way.

The debugging facilities can be used for other purposes. It is easy to modify the system to allow assertions along the code, for instance to check loop invariants.

For the behaviour of the system is case of failure of the debugging assertions we rely on the underlying debug facilities (the effect of `assert` in C++ in our case). Of course it is possible to modify this behaviour to use failures to locate more precisely the bug and the wrong function in the vein of algorithmic debugging [16]. More precisely, when a function call violates its postcondition, function invocations inside the code are checked. Those that have an incomplete postcondition (i.e. annotated as **and_check**or **either_check**) are on suspicion of wrong behaviour.

Another relatively easy extension is to use the SLAM debugging facilities for program validation by testing automatically sequences of data designed to cover all the cases.

References

[1] David Abramson, Ian Foster, John Michalakes, and Rok Socič. Relative debugging: A new methodology for debugging scientific applications. *Communications of the ACM*, 39(11):69–77, November 1996. Uses IBM Data Explorer to illustrate the computational differences between two versions of the same program, using the Guard (Griffith University Relative Debugger) system to automatically carry out run-time monitoring and comparison of results of two implementations of the same program, possibly written in different programming languages, and possibly running on different computer architectures at different geographic locations.

[2] Manuel Blum and Hal Wasserman. Software reliability via run-time result-checking. *Journal of the ACM*, 44(6):826–849, November 1997.

[3] Manuel Clavel, Francisco Durán, Steven Eker, Patrick Lincoln, Narciso Martí-Oliet, José Meseguer, and José Quesada. *A Maude Tutorial*. CSL, SRI International, March 2000.

[4] Joseph A. Goguen, Timothy Winkler, Jose Meseguer, Kokichi Futatsugi, and Jean-Pierre Jouannaud. Introducing OBJ. Technical report, Oxford + SRI, October 1993.

[5] John V. Guttag and James J. Horning. A tutorial on Larch and LCL, a Larch/C interface language. In S. Prehn and W. J. Toetenel, editors, *VDM91: Formal Soft-*

ware Development Methods. Springer-Verlag Lecture Notes in Computer Science 551, October 1991.

[6] M. Hermenegildo, F. Bueno, D. Cabeza, M. Carro, M. García de la Banda, P. López-García, and G. Puebla. The Ciao Multi-Dialect Compiler and System: An Experimentation Workbench for Future (C)LP Systems. In *Parallelism and Implementation of Logic and Constraint Logic Programming*, pages 65–85. Nova Science, Commack, NY, USA, April 1999.

[7] M. Hermenegildo, F. Bueno, G. Puebla, and P. López. Program Analysis, Debugging and Optimization Using the Ciao System Preprocessor. In *1999 International Conference on Logic Programming*, Cambridge, MA, November 1999. MIT Press.

[8] P. Hudak, J. Fasel, and J. Peterson. A gentle introduction to haskell. Technical Report YALEU/DCS/RR-90, Yale University, 1996.

[9] Cliff B. Jones. *Systematic Software Development Using VDM*. Prentice Hall, 1986.

[10] S. Peyton Jones and J. Hughes, editors. *Report on the Programming Language Haskell 98. A non-strict Purely Functional Language*. February 1999.

[11] David Luckham. *Programming with Specifications: An Introduction to Anna, A Language for Specifying Ada Programs*. Texts and Monographs in Computer Science. Springer-Verlag, New York, NY, 1990.

[12] Martin Odersky and Philip Wadler. Pizza into Java: Translating theory into practice. In *Proc. 24th ACM Symposium on Principles of Programming Languages*, January 1997.

[13] G. Puebla. *Advanced Compilation Techniques based on Abstract Interpretation and Program transformation*. PhD thesis, Universidad Politécnica de Madrid (UPM), Facultad Informatica UPM, 28660-Boadilla del Monte, Madrid-Spain, November 1997.

[14] G. Puebla, F. Bueno, and M. Hermenegildo. A Framework for Assertion-based Debugging in Constraint Logic Programming. In *Proceedings of the JICSLP'98 Workshop on Types for CLP*, Manchester, UK, June 1998.

[15] L. Rapanotti and A. Socorro. Introducing FOOPS. Technical Report PRG-TR-28-92, Oxford University Computing Laboratory, 1992.

[16] E.Y. Shapiro. *Algorithmic Program Debugging*. The MIT Press, 1982.

[17] The SLAM site. `http://lml.ls.fi.upm.es/slam`. Includes a draft of the language definition.

[18] B. Stroustrup. *The Programming Language C++*. Addison Wesley, 1990.

[19] J. B. Wordsworth. *Software Development with Z*. Addison-Wesley, 1992.